Polyhedral Bounds for Forbidden-Vertices Sets and No-Good Cut Relaxations

Gustavo Angulo*

July 1, 2026

Abstract

We study the convex hull obtained after deleting prescribed vertices from the binary cube. The analysis separates three regimes according to the number of deleted vertices. When this number is fixed, both the original-space facet count and the linear extension complexity remain linear in the ambient dimension, up to constants depending only on the number of deleted vertices. When the deleted set grows, this behavior disappears: polynomially many forbidden vertices can force superpolynomially many facets and superlinear extension complexity, while unrestricted deletions can reach superexponentially many facets and exponential extension complexity. We also compare the no-good relaxation with known valid inequalities for structured forbidden sets, deriving exact violation bounds and a worst-case gap after adding all inequalities considered.

1 Introduction

Binary optimization models often encode admissible decisions as vertices of the unit cube. In decomposition algorithms with binary master variables, the most basic way to exclude a previously evaluated assignment $\bar{x} \in \{0,1\}^n$ is the no-good inequality

$$\sum_{i:\bar{x}_i=0} x_i + \sum_{i:\bar{x}_i=1} (1 - x_i) \geq 1.$$

Such inequalities appear naturally in Benders-type methods [4], including the integer L-shaped method [16] and logic-based Benders decomposition [11]. Their appeal is their universality: they require only the binary incumbent and immediately prevent the master problem from returning the same assignment.

This universality is also the source of their weakness. A no-good cut is a point exclusion. Except at the incumbent, it carries little information about the recourse function, the subproblem feasible region, or the geometry of the admissible binary points. In integer L-shaped algorithms, no-good optimality cuts therefore provide a finite but local approximation of the value function. In logic-based [7] and combinatorial Benders methods [5], the same limitation motivates stronger cuts derived from subproblem inference rather than from the incumbent alone.

---

[*] gangulo@uc.cl, Department of Industrial and Systems Engineering, Pontificia Universidad Católica de Chile, Santiago, Chile 7820436.

A common strengthening is to reduce the support of the no-good cut. If infeasibility or a bound on the recourse value is certified by a subset $S \subseteq [n]$ of the master decisions, then the cut

$$\sum_{i \in S : \bar{x}_i = 0} x_i + \sum_{i \in S : \bar{x}_i = 1} (1 - x_i) \geq 1$$

excludes every assignment that repeats the same pattern on S. This principle underlies combinatorial Benders cuts and related IIS/MIS-based cuts: a small certificate can yield a substantially stronger master cut. The tradeoff is that finding such a certificate is problem dependent and may require solving an auxiliary combinatorial problem [5, 12].

The question of describing a set of binary points in the original space by linear inequalities goes back to [14], where it is shown that every subset S of the vertices of the n-cube can be defined by at most 2^{n-1} linear inequalities, and that every value $1 \leq k \leq 2^{n-1}$ is attained as the minimum number of such inequalities for some S, already by canonical (no-good) cuts. Original-space convexification of binary sets has also been studied through Chvátal–Gomory rank. In [6], it is shown that certain polytopes contained in the binary cube have small Chvátal rank when the subgraph induced by the missing vertices has very restricted treewidth, while [3] generalizes this line of work using the parameters notch and gap, obtaining bounded Chvátal–Gomory rank under broader structural conditions. These results show that the structure of the missing vertices is central to cutting-plane convexification.

A complementary approach is to allow auxiliary variables and work in an extended space, where the relevant measure is the extension complexity introduced in [18]. The work [1] introduces the forbidden-vertices problem for a general polytope and proves tractability and extended-formulation results when the underlying polytope has binary vertices, showing that forbidden vertices can often be handled compactly after lifting. Such results do not, however, determine the number of facets needed in the original space, and the relative power of the two viewpoints is itself well studied: [15] investigates the least number of inequalities of any formulation in the original variables and proves that auxiliary variables are essential in many cases, giving exponential lower bounds on the number of inequalities needed without them. The extended space is not always cheap either, since general binary polytopes may require exponentially large extended formulations [17]; and already in the original space, the maximal number of facets of a 0/1 polytope is superexponential in the dimension [2]. Closest to our setting, [13] shows that every $A \subseteq \{0,1\}^n$ admits a polytope P with $P \cap \{0,1\}^n = A$ and extension complexity $O(2^{n/2})$, whereas some A force extension complexity $2^{n(1-o(1))/3}$ for every such P; we instead require P to be exactly the convex hull of the retained vertices, which is a more restrictive demand. This leaves open the quantitative behavior of forbidden-vertices polytopes as a function of the ambient dimension and the number of forbidden vertices.

The discussion so far concerns exact descriptions, but in practice one rarely adds the full convex hull: the cheapest reaction to a forbidden vertex is to append its single no-good cut, giving a relaxation that lies between the cube $[0,1]^n$ and the exact convex hull of the retained vertices. Since the quantity that controls bound quality in branch-and-cut and in Benders-type methods is the optimal value of a linear objective over the current relaxation, a natural way to measure how much is lost is to compare the optimum of a nonnegative objective over the no-good relaxation with its optimum over the exact convex hull. The same question can be asked of the stronger cuts proposed for these sets: [6] introduced families of valid inequalities that exploit the local structure of the forbidden set, and it is natural to ask how much each of them recovers and whether, taken together, they close the gap to the exact hull.

1.1 Notation and definitions

We write $[n] := \{1, \ldots, n\}$. Given a set $X \subseteq \{0,1\}^n$ of forbidden vertices, the central object is

$$P_n(X) := \operatorname{conv}(\{0,1\}^n \setminus X).$$

We let $k := |X|$ and refer to the elements of $\{0,1\}^n \setminus X$ as feasible vertices. The expression $v \oplus e^i$ denotes the vertex obtained from v by flipping coordinate i, and $N_X(v) := \{i \in [n] : v \oplus e^i \in X\}$ denotes the set of forbidden-neighbor directions of v. A facet of $P_n(X)$ is called a cube facet if it is induced by one of the bounds $0 \leq x_i \leq 1$; all other facets are called non-cube facets.

For a forbidden vertex $v \in X$, the affine Hamming distance from x to v is

$$d_v(x) := \sum_{j : v_j = 0} x_j + \sum_{j : v_j = 1} (1 - x_j).$$

The no-good relaxation associated with X is $Q_n(X) := \{x \in [0,1]^n : d_v(x) \geq 1 \text{ for all } v \in X\}$. Thus $P_n(X) = \operatorname{conv}(Q_n(X) \cap \{0,1\}^n) \subseteq Q_n(X)$. For $|X| = 1$, the cube bounds together with the corresponding no-good inequality give the exact convex hull. For larger X, interactions among forbidden vertices can create facets that are not obtained by simply adding one no-good inequality per forbidden point.

We measure original-space complexity by the number $f(P)$ of facets of a polytope P, and lifted complexity by its linear extension complexity $\operatorname{xc}(P)$. For admissible pairs (n,k), define

$$F(n,k) := \max\{f(P_n(X)) : \ X \subseteq \{0,1\}^n, \ |X| \leq k\}$$

and

$$\operatorname{XC}(n,k) := \max\{\operatorname{xc}(P_n(X)) : \ X \subseteq \{0,1\}^n, \ |X| \leq k\}.$$

Here admissible means $1 \leq k < 2^n$. We distinguish three regimes: fixed k, where k is independent of n; polynomial k, where k grows as a fixed polynomial in the ambient dimension; and arbitrary k, where no growth restriction is imposed.

1.2 Contributions and organization

The contributions are as follows; see also Table 1.

1. **Facets:** For fixed k, we give a normal form for non-cube facets and count the possible coefficient patterns. This yields $F(n,k) \leq 2n + O_k(1)$, while a family of independent forbidden weight-two vertices gives the matching lower bound $F(n,k) \geq 2n + k$ for all sufficiently large n. We also show that polynomially many forbidden vertices can already force superpolynomially many facets by embedding hard lower-dimensional binary polytopes as faces. In the unrestricted regime, known extremal bounds for binary polytopes give the sharp qualitative scale $2^{\Theta(n \log n)}$ [10].

2. **Extension complexity:** Fixed k again gives $\operatorname{XC}(n,k) = 2n + O_k(1)$. In contrast, polynomially many forbidden vertices can force superlinear extension complexity, even though a general lifted formulation of size $O(nk)$ is available. If k is unrestricted, face embeddings give exponential extension complexity [17].

3. **No-good relaxation:** We quantify the gap between the no-good relaxation and known structured cuts [6]. We prove a square-free dichotomy for multiplicative gaps, exact violation bounds for path, subcube, star, tulip, and propeller inequalities, and a worst-case $\Theta(n)$ gap after adding all inequalities considered in the paper.

Table 1: Summary of the bounds proved or used in this paper. Constants hidden in $O(\cdot)$, $\Omega(\cdot)$, and $o(\cdot)$ may depend on fixed exponents.

Regime for k	Facets	Extension complexity
Fixed k	$F(n,k) = 2n + O_k(1)$, and the dependence on n is tight up to an additive constant depending only on k.	$\mathrm{XC}(n,k) = 2n + O_k(1)$.
Polynomial k	For every fixed $\rho > 0$, there are sets with $k \le n^\rho$ and $f(P_n(X)) \ge n^{\Omega_\rho(\log\log n)}$. Hence no polynomial in (n,k) bounds the original-space facet count.	There is a general lifted upper bound $\mathrm{xc}(P_n(X)) \le O(nk)$. For every fixed $\rho > 0$, some sets with $k \le n^\rho$ satisfy $\mathrm{xc}(P_n(X)) \ge n^{\rho/2-o(1)}$.
Arbitrary k	For all k, $F(n,k) \le 2^{O(n\log n)}$; for some k_n, $F(n,k_n) \ge 2^{\Omega(n\log n)}$.	The upper bound $\mathrm{xc}(P_n(X)) \le O(nk)$ holds for all X; for some sets, $\mathrm{xc}(P_n(X)) \ge 2^{n/2-o(n)}$.

The rest of the paper is organized as follows. Section 2 proves the facet bounds for fixed, polynomial, and arbitrary k. Then, Section 3 proves the extension-complexity bounds in the same three regimes. Section 4 studies gaps of no-good cuts and the known structured inequalities considered here. Finally, Section 5 draws some concluding remarks.

2 Facet bounds

This section studies $F(n,k)$ in the three regimes described in Table 1. The fixed-k analysis is original-space and constructive: it separates cube facets from all remaining facets. The growing-k lower bounds are obtained by embedding hard 0/1 polytopes as faces.

2.1 Fixed number of forbidden vertices

We first prove that, for fixed $k = |X|$, the number of facets of $P_n(X)$ is linear in n. The proof separates the possible cube facets from the remaining facets. The key observation is that every non-cube facet can have nontrivial coefficients only in directions pointing from a suitably chosen forbidden vertex to other forbidden vertices.

Lemma 1. *Suppose that $P_n(X)$ is full-dimensional. Let $\alpha^\top x \le \beta$ define a facet of $P_n(X)$ that is not a facet of the cube $[0,1]^n$. Then there exist a forbidden vertex $v \in X$, a set $D \subseteq N_X(v)$, and coefficients $c_i \in [0,1)$ for $i \in D$, such that, after the change of coordinates*

$$y_i = \begin{cases} x_i, & v_i = 0, \\ 1 - x_i, & v_i = 1, \end{cases}$$

the facet-defining inequality can be written as

$$\sum_{i \notin D} y_i + \sum_{i \in D} c_i y_i \geq 1.$$

In particular, $|D| \leq |X| - 1$.

Proof. The inequality cannot be valid for the whole cube. Indeed, if it were valid for $[0,1]^n$, then the cube face exposed by the same hyperplane would contain the $(n-1)$-dimensional face of $P_n(X)$ defined by $\alpha^\top x = \beta$. Hence the exposed cube face would have dimension at least $n-1$, and the inequality would be a cube facet, contrary to assumption.

Thus the inequality cuts off at least one cube vertex. Since it is valid for $P_n(X)$, every cube vertex that violates it is forbidden. Let $Y := \{u \in \{0,1\}^n : \alpha^\top u > \beta\}$. Then $\emptyset \neq Y \subseteq X$. Choose $v \in Y$ maximizing $\alpha^\top u$ over X, and set $\delta := \alpha^\top v - \beta > 0$.

Apply the coordinate transformation above, which sends v to the origin. In the y-coordinates, the original inequality is equivalent to $\sum_{i=1}^n a_i y_i \geq 1$, where

$$a_i := \frac{\alpha^\top v - \alpha^\top (v \oplus e^i)}{\delta}.$$

We first show that $a_i \geq 0$ for all i. If $v \oplus e^i \in X$, this follows from the choice of v as a maximizer of $\alpha^\top u$ over X. If $v \oplus e^i \notin X$, then $v \oplus e^i$ is feasible for $P_n(X)$, so $\alpha^\top(v \oplus e^i) \leq \beta < \alpha^\top v$, and hence $a_i > 0$.

If $a_i < 1$, then the binary point e^i in the y-coordinates violates $\sum_j a_j y_j \geq 1$. Hence the corresponding cube vertex $v \oplus e^i$ must be forbidden. Therefore $D := \{i \in [n] : a_i < 1\}$ is contained in $N_X(v)$.

It remains to show that $a_i = 1$ for all $i \notin D$. By definition, $a_i \geq 1$ for $i \notin D$. Suppose that $a_i > 1$ for some $i \notin D$. Since all coefficients are nonnegative, no point satisfying $\sum_j a_j y_j = 1$ can have $y_i = 1$. Thus the exposed face is contained both in $\sum_j a_j y_j = 1$ and in $y_i = 0$. These two hyperplanes are distinct, so the face has dimension at most $n-2$, contradicting that the inequality defines a facet. Therefore $a_i = 1$ for all $i \notin D$.

The claimed representation follows. Finally, $D \subseteq N_X(v)$, and the map $i \mapsto v \oplus e^i$ is injective from $N_X(v)$ into $X \setminus \{v\}$. Hence $|D| \leq |N_X(v)| \leq |X| - 1$. ☐

Lemma 2. *Suppose that $P_n(X)$ is full-dimensional. Fix $v \in X$ and $D \subseteq N_X(v)$, and let $q := |D|$. There are at most $2^{q(q+1)}$ facet-defining inequalities of $P_n(X)$ which, in the coordinates centered at v, have the form*

$$\sum_{i \notin D} y_i + \sum_{i \in D} c_i y_i \geq 1, \qquad 0 \leq c_i < 1.$$

Proof. For fixed v and D, only the vector $c = (c_i)_{i \in D} \in \mathbb{R}^q$ is unknown. Consider one facet of the stated form, and let T be the set of vertices of $P_n(X)$ satisfying it at equality. Since the inequality defines a facet, $\operatorname{conv}(T)$ has dimension $n-1$.

The projections of the points in T onto the coordinates in D must linearly span $\mathbb{R}^q$. Otherwise, there is a nonzero $h \in \mathbb{R}^q$ such that $h^\top y_D = 0$ for all $y \in T$. The facet would then be contained in two distinct hyperplanes, namely the facet hyperplane and $h^\top y_D = 0$, and hence would have dimension at most $n-2$, a contradiction.

Thus there exist tight vertices whose projections $z^1, \dots, z^q \in \{0,1\}^q$ are linearly independent. For each such vertex, the equality condition gives $\sum_{i \notin D} y_i + c^\top z^r = 1$. The first term is a nonnegative integer and $c^\top z^r \geq 0$, so $c^\top z^r \in \{0,1\}$ for every r. Hence c is determined by a system $Mc = b$, where M is an invertible binary $q \times q$ matrix and $b \in \{0,1\}^q$. There are at most 2^{q^2} choices for M and at most 2^q choices for b, giving at most $2^{q(q+1)}$ possible vectors c. $\square$

Theorem 3. *We have*

$$f(P_n(X)) \leq 2n + k \sum_{q=0}^{k-1} \binom{k-1}{q} 2^{q(q+1)}.$$

In particular,

$$f(P_n(X)) \leq 2n + k 2^{k^2-1}.$$

Proof. First, note that every proper affine hyperplane contains at most 2^{n-1} vertices of $\{0,1\}^n$. Indeed, if the hyperplane is given by $a^\top x = \beta$ with $a_j \neq 0$, then, after fixing the $n-1$ coordinates different from j, there is at most one value of x_j satisfying the equation.

Suppose first that $k < 2^{n-1}$. Then $|\{0,1\}^n \setminus X| = 2^n - k > 2^{n-1}$, so $\{0,1\}^n \setminus X$ is not contained in any proper affine hyperplane. Hence $P_n(X)$ is full-dimensional. In this case there are at most $2n$ cube facets. Every remaining facet is associated, by Lemma 1, with a forbidden vertex $v \in X$ and a set $D \subseteq N_X(v)$. Since $N_X(v) \subseteq X \setminus \{v\}$, for fixed v there are at most $\binom{k-1}{q}$ possible choices of D with $|D| = q$. By Lemma 2, each such pair (v, D) gives at most $2^{q(q+1)}$ possible inequalities. Therefore

$$f(P_n(X)) \leq 2n + k \sum_{q=0}^{k-1} \binom{k-1}{q} 2^{q(q+1)}.$$

It remains to consider the case $k \geq 2^{n-1}$. Let $r := |\{0,1\}^n \setminus X| = 2^n - k$. Then $r \leq k$. A polytope with r vertices has at most 2^r faces, and hence at most 2^r facets. Therefore

$$f(P_n(X)) \leq 2^r \leq 2^k.$$

If $k = 1$, then necessarily $n = 1$ and $P_n(X)$ is a single point, so $f(P_n(X)) = 0$ and the desired bound holds. If $k \geq 2$, then the term corresponding to $q = k-1$ gives

$$\sum_{q=0}^{k-1} \binom{k-1}{q} 2^{q(q+1)} \geq 2^{(k-1)k} \geq 2^k.$$

Thus the desired bound also holds when $k \geq 2^{n-1}$.

Finally, since $q(q+1) \leq (k-1)k$ for $0 \leq q \leq k-1$ and $\sum_{q=0}^{k-1} \binom{k-1}{q} = 2^{k-1}$,

$$\sum_{q=0}^{k-1} \binom{k-1}{q} 2^{q(q+1)} \leq 2^{(k-1)k} \sum_{q=0}^{k-1} \binom{k-1}{q} = 2^{k^2-1}.$$

This proves the simplified estimate. $\square$

The preceding upper bound is tight in its dependence on n when k is fixed. We use Corollary 1 of [1] in the special case $P = [0,1]^n$. Namely, if $X \subseteq \{0,1\}^n$ is an independent set in the graph of the n-cube, then the convex hull obtained by deleting the vertices in X is described by the cube inequalities together with one neighboring halfspace for each deleted vertex:

$$P_n(X) = \operatorname{conv}(\{0,1\}^n \setminus X) = [0,1]^n \cap \bigcap_{v \in X} H_v,$$

where $H_v := \{x \in \mathbb{R}^n : d_v(x) \geq 1\}$.

Theorem 4. *For every fixed k and all sufficiently large n, $F(n,k) \geq 2n + k$. Therefore, $F(n,k) = 2n + O_k(1) = \Theta_k(n)$.*

Proof. Let n be large enough that $\binom{n}{2} \geq k$ and $n \geq 4$. Choose X as any set of k distinct vertices of Hamming weight two. No two vertices of X are adjacent in the cube, and all cube vertices are simple. Therefore, Corollary 1 of [1] gives $P_n(X) = Q_n(X)$.

For each $v \in X$, all n neighbors of v have Hamming weight one or three, and hence none belongs to X. These neighbors are affinely independent in the hyperplane of the no-good cut, so this cut defines a facet of $P_n(X)$. This contributes k non-cube facets.

The $2n$ cube facets also survive. For $x_i = 0$, the vertices 0 and e^j with $j \neq i$ are feasible and affinely independent in the hyperplane $x_i = 0$. For $x_i = 1$, the vertices $\mathbf{1}$ and $\mathbf{1} - e^j$ with $j \neq i$ are feasible for $n \geq 4$ and affinely independent in the hyperplane $x_i = 1$. Hence each cube facet remains a facet of $P_n(X)$.

Thus $P_n(X)$ has at least $2n + k$ facets. Combining this with Theorem 3 gives $F(n,k) = 2n + O_k(1) = \Theta_k(n)$. □

2.2 Polynomially many forbidden vertices

We next show that polynomially many forbidden vertices already suffice to force a superpolynomial number of facets in the original space. The construction embeds a hard lower-dimensional 0/1-polytope as a face of a larger forbidden-vertices polytope.

We shall use the following results.

Theorem 5 ([10])**.** *There exists an absolute constant $c > 0$ such that, for every sufficiently large integer d, there is a full-dimensional 0/1-polytope $Q_d = \operatorname{conv}(S_d) \subseteq \mathbb{R}^d$, where $S_d \subseteq \{0,1\}^d$, satisfying*

$$f(Q_d) \geq \left(\frac{cd}{\log d}\right)^{d/2}.$$

Lemma 6 (see for instance [19, Theorem 2.7])**.** *Let P be a polytope and let F be a nonempty face of P. Then $f(F) \leq f(P)$.*

Theorem 7. *Fix $\rho > 0$. Then there exists a sequence $X_n \subseteq \{0,1\}^n$ such that $|X_n| \leq n^\rho$ and*

$$f(P_n(X_n)) \geq \left(\frac{c_\rho \log n}{\log\log n}\right)^{\frac{\rho}{2}\log_2 n - O(1)} = n^{\Omega_\rho(\log\log n)}.$$

In particular, even when $|X_n|$ is polynomial in n, the number of facets of $P_n(X_n)$ is not bounded by any polynomial in $(n, |X_n|)$.

Proof. Let $k_n := \lfloor n^\rho \rfloor$ and set $d_n := \lfloor \log_2 k_n \rfloor - 2$. Then $d_n = \rho \log_2 n - O(1)$ and $2^{d_n} \le k_n/4$.

Apply Theorem 5 with $d = d_n$. Thus, for all sufficiently large n, there is a 0/1-polytope $Q_{d_n} = \operatorname{conv}(S_{d_n}) \subseteq \mathbb{R}^{d_n}$, with $S_{d_n} \subseteq \{0,1\}^{d_n}$, such that

$$f(Q_{d_n}) \ge \left(\frac{c d_n}{\log d_n} \right)^{d_n/2}$$

for some absolute constant $c > 0$.

Write $r_n := n - d_n$ and decompose $\{0,1\}^n = \{0,1\}^{d_n} \times \{0,1\}^{r_n}$. Let $G_n := \{(u,z) \in \{0,1\}^{d_n} \times \{0,1\}^{r_n} : z = 0\}$ be the coordinate face where the last r_n coordinates are zero. Inside this face, forbid exactly the complement of S_{d_n}, thus $X_n := \big(\{0,1\}^{d_n} \setminus S_{d_n}\big) \times \{0\}^{r_n}$. Then $|X_n| \le 2^{d_n} \le k_n/4 \le \lfloor n^\rho \rfloor$.

The face of $P_n(X_n)$ exposed by $\sum_{j=1}^{r_n} z_j = 0$ is

$$P_n(X_n) \cap \{z = 0\} = \operatorname{conv}(S_{d_n}) \times \{0\}^{r_n} = Q_{d_n} \times \{0\}^{r_n}.$$

Therefore Q_{d_n} is affinely isomorphic to a face of $P_n(X_n)$. By Lemma 6, $f(P_n(X_n)) \ge f(Q_{d_n})$. Hence

$$f(P_n(X_n)) \ge \left(\frac{c d_n}{\log d_n} \right)^{d_n/2}.$$

Since $d_n = \rho \log_2 n - O(1)$, there exists a constant $c_\rho > 0$ such that

$$f(P_n(X_n)) \ge \left(\frac{c_\rho \log n}{\log \log n} \right)^{\frac{\rho}{2} \log_2 n - O(1)} = n^{\Omega_\rho(\log \log n)}.$$

The last quantity is superpolynomial in n. Since $|X_n| \le n^\rho$, this also rules out any polynomial upper bound in the natural parameters $(n, |X_n|)$. □

2.3 Arbitrary number of forbidden vertices

The fixed-k regime is linear, and the previous subsection shows that the polynomial regime is already superpolynomial in the original space. We now derive general upper bounds and a matching qualitative lower bound for unrestricted k.

Theorem 8. *For all admissible (n,k) with sufficiently large n,*

$$F(n,k) \le \min\{2n + k 2^{k^2-1},\ 2^{O(n \log n)}\}.$$

Proof. The first bound is Theorem 3. For the second bound, let

$$\phi(\delta) := \max\{f(Q) : Q = \operatorname{conv}(S),\ S \subseteq \{0,1\}^m \text{ for some } m \ge \delta,\ \dim Q = \delta\}$$

be the maximum number of facets of a δ-dimensional 0/1-polytope, taken over all ambient dimensions. We first note that ϕ is non-decreasing. Indeed, if Q is a δ-dimensional 0/1-polytope attaining $\phi(\delta)$, then the prism $Q \times [0,1]$ is a $(\delta+1)$-dimensional 0/1-polytope with $f(Q \times [0,1]) = f(Q) + 2$, so $\phi(\delta+1) \ge \phi(\delta) + 2$.

Now $P_n(X)$ is a 0/1-polytope of dimension $d := \dim P_n(X) \le n$, so $f(P_n(X)) \le \phi(d) \le \phi(n)$. By [9] there is an absolute constant $C > 0$ with $\phi(n) \le C(n-2)!$ for all sufficiently large n. Hence $f(P_n(X)) \le C(n-2)!$, and Stirling's formula gives $C(n-2)! = 2^{O(n \log n)}$. □

The next theorem shows that the factorial-type bound is qualitatively tight in n: forbidden-vertices polytopes can have $2^{\Omega(n \log n)}$ facets.

Theorem 9. *There exist a constant $c > 0$ and a sequence $k_n \leq 2^n$ such that*

$$F(n, k_n) \geq \left(\frac{cn}{\log n}\right)^{n/2}.$$

In particular, $F(n, k_n) = 2^{\Omega(n \log n)}$. Moreover, along this sequence, $F(n, k_n) \geq k_n^{\Omega(\log \log k_n)}$.

Proof. Apply Theorem 5 with $d = n$. Thus, for all sufficiently large n, there is a full-dimensional 0/1-polytope $Q_n = \operatorname{conv}(S_n) \subseteq \mathbb{R}^n$, with $S_n \subseteq \{0,1\}^n$, such that

$$f(Q_n) \geq \left(\frac{cn}{\log n}\right)^{n/2}.$$

Define $X_n := \{0,1\}^n \setminus S_n$ and $k_n := |X_n|$. Then $k_n \leq 2^n$ and $P_n(X_n) = \operatorname{conv}(S_n) = Q_n$. Therefore

$$F(n, k_n) \geq f(P_n(X_n)) = f(Q_n) \geq \left(\frac{cn}{\log n}\right)^{n/2}.$$

The right-hand side is $2^{\Omega(n \log n)}$.

It remains to express this lower bound in terms of k_n. The preceding bound implies that k_n is unbounded; otherwise, Theorem 3 would give $F(n, k_n) = O(n)$ along an infinite subsequence, contradicting $F(n, k_n) \geq 2^{\Omega(n \log n)}$. Hence, for all sufficiently large n, we may assume $k_n \geq 3$.

Since $k_n \leq 2^n$, we have $1 < \log_2 k_n \leq n$ and $0 < \log_2 \log_2 k_n \leq \log_2 n$. Hence

$$0 < (\log_2 k_n)(\log_2 \log_2 k_n) \leq n \log_2 n.$$

Thus

$$F(n, k_n) \geq 2^{\Omega(n \log n)} \geq 2^{\Omega((\log_2 k_n)(\log_2 \log_2 k_n))} = k_n^{\Omega(\log \log k_n)}.$$

□

3 Extension-complexity bounds

This section studies $\mathrm{XC}(n, k)$. We first recall a general lifted upper bound for forbidden vertices in the cube. We then prove that fixed k is essentially as easy as the cube, while polynomial and unrestricted k can force substantially larger extension complexity.

Proposition 10 ([1])**.** *For every nonempty $X \subseteq \{0,1\}^n$, the forbidden-vertices polytope satisfies* $\mathrm{xc}(P_n(X)) \leq O(n|X|)$.

We shall use the following result.

Lemma 11 (see for instance [8, Remark 2.7])**.** *Let P be a polytope and let F be a face of P. Then* $\mathrm{xc}(F) \leq \mathrm{xc}(P)$.

3.1 Fixed number of forbidden vertices

We now derive the extension-complexity analogue of Theorem 3. For fixed k, the extension complexity of $P_n(X)$ is asymptotically the same as that of the cube, up to an additive term depending only on k.

Theorem 12. *Set* $t_k := \lceil \log_2(k+1) \rceil$. *Then*

$$2(n - t_k) \leq \mathrm{xc}(P_n(X)) \leq 2n + k \sum_{q=0}^{k-1} \binom{k-1}{q} 2^{q(q+1)} \leq 2n + k2^{k^2-1}.$$

Therefore, for every fixed k, $\mathrm{XC}(n,k) = 2n + O_k(1)$.

Proof. The upper bound is immediate from the original-space facet bound. Indeed, every linear description in the original space is an extended formulation, and hence $\mathrm{xc}(P_n(X)) \leq f(P_n(X))$. Theorem 3 gives

$$f(P_n(X)) \leq 2n + k \sum_{q=0}^{k-1} \binom{k-1}{q} 2^{q(q+1)} \leq 2n + k2^{k^2-1}.$$

It remains to prove the lower bound. Let $T \subseteq [n]$ be any set with $|T| = t_k$. Since $2^{t_k} \geq k+1$, there is a pattern $a \in \{0,1\}^T$ such that $v_T \neq a$ for each $v \in X$. Thus the coordinate face

$$C := \{x \in [0,1]^n : x_i = a_i \text{ for all } i \in T\}$$

contains no forbidden vertex. Therefore all binary vertices of C remain feasible for $P_n(X)$, and $P_n(X) \cap C = C$. In particular, C is a face of $P_n(X)$ affinely isomorphic to the $(n-t_k)$-dimensional cube $[0,1]^{n-t_k}$.

By Lemma 11, $\mathrm{xc}(P_n(X)) \geq \mathrm{xc}(C)$. Since C is an $(n-t_k)$-dimensional cube and the extension complexity of a d-dimensional cube is $2d$, we obtain $\mathrm{xc}(P_n(X)) \geq 2(n - t_k)$.

Combining the lower and upper bounds yields

$$2n - O_k(1) \leq \mathrm{xc}(P_n(X)) \leq 2n + O_k(1).$$

Taking the maximum over all sets $X \subseteq \{0,1\}^n$ with $|X| \leq k$ gives $\mathrm{XC}(n,k) = 2n + O_k(1)$ for every fixed k. □

3.2 Polynomially many forbidden vertices

We now show that polynomially many forbidden vertices already suffice to force superlinear extension complexity. The construction embeds, as a face of a larger forbidden-vertices polytope, a hard 0/1-polytope of logarithmic dimension.

We shall use the following result.

Theorem 13 ([17])**.** *For every sufficiently large integer* d, *there exists a set* $S_d \subseteq \{0,1\}^d$ *such that, for* $Q_d := \mathrm{conv}(S_d) \subseteq \mathbb{R}^d$,

$$\mathrm{xc}(Q_d) \geq 2^{d/2 - o(d)}.$$

Theorem 14. *Fix $\rho > 0$. Then there exists a sequence $X_n \subseteq \{0,1\}^n$ such that $|X_n| \le n^\rho$ and*

$$\mathrm{xc}(P_n(X_n)) \ge n^{\rho/2-o(1)}.$$

In particular, the lower bound is superlinear in n whenever $\rho > 2$.

Proof. Let $k_n := \lfloor n^\rho \rfloor$ and choose $d_n := \lfloor \log_2 k_n \rfloor - 2$. Then $d_n = \rho \log_2 n - O(1)$ and $2^{d_n} \le k_n/4$.

Apply Theorem 13 with $d = d_n$. Thus, for all sufficiently large n, there exists a set $S_{d_n} \subseteq \{0,1\}^{d_n}$ such that $Q_{d_n} := \mathrm{conv}(S_{d_n})$ satisfies

$$\mathrm{xc}(Q_{d_n}) \ge 2^{d_n/2-o(d_n)}.$$

Write $r_n := n - d_n$ and decompose $\{0,1\}^n = \{0,1\}^{d_n} \times \{0,1\}^{r_n}$. Let $G_n := \{(u,z) \in \{0,1\}^{d_n} \times \{0,1\}^{r_n} : z = 0\}$ be the coordinate face where the last r_n coordinates are zero. Inside this face, forbid precisely the vertices corresponding to $\{0,1\}^{d_n} \setminus S_{d_n}$. Thus $X_n := \left(\{0,1\}^{d_n} \setminus S_{d_n}\right) \times \{0\}^{r_n}$ and $|X_n| \le 2^{d_n} \le k_n/4$.

The face of $P_n(X_n)$ exposed by $\sum_{j=1}^{r_n} z_j = 0$ is

$$P_n(X_n) \cap \{z = 0\} = \mathrm{conv}(S_{d_n}) \times \{0\}^{r_n} = Q_{d_n} \times \{0\}^{r_n}.$$

Thus Q_{d_n} is affinely isomorphic to a face of $P_n(X_n)$. By Lemma 11,

$$\mathrm{xc}(P_n(X_n)) \ge \mathrm{xc}(Q_{d_n}) \ge 2^{d_n/2-o(d_n)}.$$

Since $d_n = \rho \log_2 n - O(1)$, we obtain

$$\mathrm{xc}(P_n(X_n)) \ge n^{\rho/2-o(1)}.$$

□

Note that if $k_n = \lfloor n^\rho \rfloor$ for a fixed $\rho > 0$, then $n^{\rho/2-o(1)} = k_n^{1/2-o(1)}$. Thus, along polynomial budgets $k_n = n^\rho$, Theorem 14 gives a lower bound of order $\sqrt{k_n}$ up to subpolynomial factors.

3.3 Arbitrary number of forbidden vertices

Theorem 14 shows that polynomially many forbidden vertices can force superlinear extension complexity. If the number of forbidden vertices is not controlled, the same face-embedding idea yields exponential extension complexity.

Theorem 15. *There exists a sequence $X_n \subseteq \{0,1\}^n$ such that*

$$\mathrm{xc}(P_n(X_n)) \ge 2^{n/2-o(n)}.$$

Proof. Let $m_n := n - 1$. By Theorem 13, there exists a set $S_{m_n} \subseteq \{0,1\}^{m_n}$ such that $Q_{m_n} := \mathrm{conv}(S_{m_n})$ satisfies

$$\mathrm{xc}(Q_{m_n}) \ge 2^{m_n/2(1-o(1))}.$$

Embed Q_{m_n} in the coordinate face of $\{0,1\}^n$ defined by the last coordinate being zero. More precisely, set $Y_n := S_{m_n} \times \{0\} \subseteq \{0,1\}^{m_n} \times \{0,1\}$ and forbid all other binary vertices, that is, $X_n := \{0,1\}^n \setminus Y_n$. Then

$$P_n(X_n) = \mathrm{conv}(Y_n) = Q_{m_n} \times \{0\}.$$

Hence $P_n(X_n)$ is affinely isomorphic to Q_{m_n}. Therefore

$$\mathrm{xc}(P_n(X_n)) = \mathrm{xc}(Q_{m_n}) \ge 2^{m_n/2(1-o(1))} = 2^{(n-1)/2(1-o(1))} = 2^{n/2-o(n)}.$$

□

4 Gaps for known cut families

This section compares the no-good relaxation with exact inequalities for several structured families of forbidden vertices. Fix $n \geq 1$, and for nonempty $X \subsetneq \{0,1\}^n$, write $P_X := P_n(X)$ and $Q_X := Q_n(X)$. We first give a general comparison for nonnegative affine objectives. Then we specialize it to paths, subcubes, stars, tulips, and propellers, which are introduced and studied in [6].

4.1 A general square-free bound

Let $\mathcal{L}_+$ denote the set of affine functions ℓ such that $\ell(x) \geq 0$ for every $x \in [0,1]^n$. We say that Q_X has multiplicative gap at most ρ with respect to P_X if

$$\min\{\ell(x) : x \in Q_X\} \geq \frac{1}{\rho} \min\{\ell(x) : x \in P_X\} \qquad \text{for every } \ell \in \mathcal{L}_+.$$

Equivalently, whenever the denominator is positive, the ratio between the value over P_X and the value over Q_X is at most ρ.

We say that X contains a square if it contains all four vertices of a two-dimensional face of the cube. Otherwise, X is called square-free.

Theorem 16. *The following statements hold.*

1. *If X contains a square, then the no-good relaxation has infinite multiplicative gap. More precisely, there exists $\ell \in \mathcal{L}_+$ such that* $\min\{\ell(x) : x \in P_X\} > 0$ *and* $\min\{\ell(x) : x \in Q_X\} = 0$.
2. *If X is square-free, then*

$$\min\{\ell(x) : x \in Q_X\} \geq \frac{1}{2} \min\{\ell(x) : x \in P_X\} \qquad \textit{for every } \ell \in \mathcal{L}_+.$$

   Thus the multiplicative gap is at most 2 and this bound is tight for $n \geq 2$.

Proof. Suppose first that X contains a square. Then there exist $I \subseteq [n]$ with $|I| = 2$, $F = [n] \setminus I$, and $a \in \{0,1\}^F$ such that

$$C := \{u \in \{0,1\}^n : u_j = a_j \text{ for all } j \in F\} \subseteq X.$$

Since $X \subsetneq \{0,1\}^n$, this square is not the whole cube, and hence $F \neq \emptyset$. Consider the affine function

$$\ell(x) := \sum_{j \in F:\ a_j = 0} x_j + \sum_{j \in F:\ a_j = 1} (1 - x_j).$$

This function is nonnegative on $[0,1]^n$. Every allowed binary vertex lies outside C, so it differs from a in at least one coordinate of F. Hence $\min\{\ell(x) : x \in P_X\} \geq 1$.

Let $\bar{x}$ be the center of the square C, defined by $\bar{x}_j = a_j$ for $j \in F$ and $\bar{x}_i = 1/2$ for $i \in I$. If $v \in C$, then $d_v(\bar{x}) = 1$. If $v \in X \setminus C$, then v differs from a in at least one coordinate of F, and therefore $d_v(\bar{x}) \geq 1$. Thus $\bar{x} \in Q_X$. However, $\ell(\bar{x}) = 0$. This proves the infinite-gap statement.

We now assume that X is square-free. Fix $\ell \in \mathcal{L}_+$. By complementing coordinates if necessary, we may assume that

$$\ell(x) = \alpha + \sum_{i \in L} w_i x_i,$$

where $\alpha \geq 0$, $L \subseteq [n]$, and $w_i > 0$ for all $i \in L$. This transformation is an automorphism of the cube and preserves the property of being square-free. Let $g(x) := \sum_{i \in L} w_i x_i$ and define

$$B := \{u \in \{0,1\}^n : u_i = 0 \text{ for all } i \in L\}.$$

Set

$$\beta := \min\{g(u) : u \in \{0,1\}^n \setminus X\}.$$

Then $\min\{\ell(x) : x \in P_X\} = \alpha + \beta$. If $\beta = 0$, then some allowed binary vertex belongs to B, and the desired inequality is immediate. Hence we assume $\beta > 0$. Then $B \subseteq X$. Since X is square-free, the face B has dimension at most one.

First suppose that B has dimension zero. After relabeling, $B = \{0\}$, $L = [n]$, and $g(x) = \sum_{i=1}^n w_i x_i$. Let $N := \{i \in [n] : e^i \in X\}$. If $i \notin N$, then e^i is allowed, so $\beta \leq w_i$. If $i, j \in N$ and $i \neq j$, then $e^i + e^j \notin X$; otherwise $\{0, e^i, e^j, e^i + e^j\} \subseteq X$ would be a forbidden square. Therefore $\beta \leq w_i + w_j$ for all distinct $i, j \in N$.

Take $x \in Q_X$. Since $0 \in X$, the no-good cut for 0 gives $\sum_i x_i \geq 1$. Moreover, for each $i \in N$, the no-good cut for e^i gives $x_i \leq \frac{1}{2} \sum_j x_j$. Let $s := \sum_{i \in N} x_i$, $z := \sum_{i \notin N} x_i$, and $T := s + z$. Thus $T \geq 1$ and $x_i \leq T/2$ for every $i \in N$.

We claim that $\sum_{i \in N} w_i x_i \geq \frac{\beta}{2}(s - z)$. Set $\lambda_i := w_i/\beta$ for $i \in N$. Then $\lambda_i + \lambda_j \geq 1$ for all distinct $i, j \in N$. If all $\lambda_i \geq 1/2$, then $\sum_{i \in N} \lambda_i x_i \geq s/2 \geq (s - z)/2$. Otherwise, choose $h \in N$ with $\lambda_h < 1/2$. For all $i \in N \setminus \{h\}$, we have $\lambda_i \geq 1 - \lambda_h$, and hence

$$\sum_{i \in N} \lambda_i x_i \geq \lambda_h x_h + (1 - \lambda_h)(s - x_h) = \lambda_h (2x_h - s) + s - x_h.$$

If $2x_h \leq s$, the right-hand side is at least $s/2$. If $2x_h > s$, it is at least $s - x_h$, and $x_h \leq T/2 = (s + z)/2$ implies $s - x_h \geq (s - z)/2$. This proves the claim.

Using also $\beta \leq w_i$ for $i \notin N$, we obtain

$$g(x) = \sum_{i \in N} w_i x_i + \sum_{i \notin N} w_i x_i \geq \frac{\beta}{2}(s - z) + \beta z = \frac{\beta}{2} T \geq \frac{\beta}{2}.$$

Now suppose that B has dimension one. After relabeling, $B = \{0, e^r\}$ for some $r \in [n]$. Then $L = [n] \setminus \{r\}$ and $g(x) = \sum_{i \neq r} w_i x_i$. Since $B \subseteq X$, both 0 and e^r are forbidden. Let $S := \sum_{i \neq r} x_i$. The no-good cuts for 0 and e^r imply $x_r + S \geq 1$ and $1 - x_r + S \geq 1$. Hence $S \geq \max\{x_r, 1 - x_r\} \geq 1/2$.

For every $i \neq r$, the two vertices e^i and $e^r + e^i$ cannot both be forbidden, because otherwise $0, e^r, e^i, e^r + e^i$ would be a forbidden square. Thus at least one of them is allowed. Both have g-value w_i, so $\beta \leq w_i$ for every $i \neq r$. It follows that

$$g(x) = \sum_{i \neq r} w_i x_i \geq \beta \sum_{i \neq r} x_i = \beta S \geq \frac{\beta}{2}.$$

In both cases, every $x \in Q_X$ satisfies $g(x) \geq \beta/2$. Therefore

$$\min\{\ell(x) : x \in Q_X\} \geq \alpha + \frac{\beta}{2} \geq \frac{\alpha + \beta}{2} = \frac{1}{2} \min\{\ell(x) : x \in P_X\},$$

where the second inequality uses $\alpha \geq 0$.

Finally, to show tightness, let $X = \{0, e^1\} \subseteq \{0,1\}^n$, with $n \geq 2$, and consider $\ell(x) = \sum_{i=2}^n x_i$. Then every allowed binary vertex has ℓ-value at least one, so $\min\{\ell(x) : x \in P_X\} = 1$. On the other hand, the two no-good cuts imply

$$\sum_{i=2}^{n} x_i \geq \max\{x_1, 1 - x_1\} \geq \frac{1}{2},$$

and equality is attained by taking $x_1 = 1/2$ and $\sum_{i=2}^n x_i = 1/2$. Hence $\min\{\ell(x) : x \in Q_X\} = 1/2$. □

4.2 Induced paths

Assume that the subgraph induced by X is a path with at least one edge and $n \geq 2$. Such a set is square-free, and therefore Theorem 16 gives a factor-two bound for every nonnegative affine objective. We now derive the corresponding sharp statement for the edge inequalities that describe the convex hull in this case.

For an edge $e = \{v, w\}$ of the path, let i be the unique coordinate in which v and w differ. Define

$$\ell_e(x) := \sum_{j \neq i} \bigl(v_j(1 - x_j) + (1 - v_j)x_j\bigr).$$

This expression is independent of the endpoint used, because $v_j = w_j$ for all $j \neq i$. For induced paths, from [6] we have

$$P_X = \{x \in [0,1]^n : \ell_e(x) \geq 1 \text{ for all } e \in E(G[X])\}.$$

Proposition 17. *For every edge $e \in E(G[X])$,*

$$\min\{\ell_e(x) : x \in Q_X\} = \frac{1}{2}.$$

Proof. The lower bound follows from Theorem 16, applied to the nonnegative affine function ℓ_e, because $\ell_e(x) \geq 1$ is valid for P_X.

To prove tightness, let $e = \{v, w\}$ and let i be the coordinate in which the endpoints differ. Choose $r \neq i$, which is possible because $n \geq 2$. Define $x^* \in [0,1]^n$ by $x^*_i = x^*_r = 1/2$ and $x^*_j = v_j$ for all $j \notin \{i, r\}$. The coordinates i and r contribute exactly $1/2$ each to $d_u(x^*)$ for every $u \in \{0,1\}^n$, so $x^* \in Q_X$. Since ℓ_e excludes coordinate i and all coordinates $j \neq i, r$ coincide with v_j, we get $\ell_e(x^*) = 1/2$. □

4.3 Forbidden subcubes

Let X induce a d-dimensional subcube. Thus there are $I \subseteq [n]$, $F = [n] \setminus I$, and $a \in \{0,1\}^F$ such that $|I| = d$ and

$$X = \{x \in \{0,1\}^n : x_j = a_j \text{ for all } j \in F\}.$$

We assume $F \neq \emptyset$, so $X \neq \{0,1\}^n$. For $J \subseteq [n]$ and $u \in \{0,1\}^J$, define

$$\Delta_J^u(x) := \sum_{j \in J:\ u_j = 0} x_j + \sum_{j \in J:\ u_j = 1} (1 - x_j).$$

The exact convex hull is described by the subcube inequality

$$P_X = \{x \in [0,1]^n : \Delta_F^a(x) \geq 1\}.$$

The no-good relaxation is

$$Q_X = \{x \in [0,1]^n : \Delta_F^a(x) + \Delta_I^z(x) \geq 1 \text{ for all } z \in \{0,1\}^I\}.$$

Proposition 18. *Let $k := |X| = 2^d$. Then*

$$\min\{\Delta_F^a(x) : x \in Q_X\} = \max\left\{0, 1 - \frac{d}{2}\right\}.$$

Proof. For fixed $x \in [0,1]^n$, minimizing the contribution of one free coordinate over $z_i \in \{0,1\}$ gives $\min\{x_i, 1 - x_i\}$. Hence the no-good family is equivalent to

$$\Delta_F^a(x) + \eta_I(x) \geq 1, \qquad \eta_I(x) := \sum_{i \in I} \min\{x_i, 1 - x_i\}.$$

Since $\eta_I(x) \leq d/2$, every $x \in Q_X$ satisfies $\Delta_F^a(x) \geq 1 - d/2$. Nonnegativity gives $\Delta_F^a(x) \geq \max\{0, 1 - d/2\}$.

The bound is attained as follows. If $d = 0$, the no-good relaxation already contains $\Delta_F^a(x) \geq 1$. If $d = 1$, let $I = \{i\}$ and choose $p \in F$. Set $x_i = x_p = 1/2$ and $x_j = a_j$ for all $j \in F \setminus \{p\}$. Then $\eta_I(x) = 1/2$ and $\Delta_F^a(x) = 1/2$. If $d \geq 2$, set $x_j = a_j$ for $j \in F$ and $x_i = 1/2$ for $i \in I$. Then $\Delta_F^a(x) = 0$ and $\eta_I(x) = d/2 \geq 1$. These cases prove the formula. $\square$

The objective-gap consequence is also immediate. For $d = 0$, the no-good and convex-hull descriptions coincide on the subcube inequality. For $d = 1$, the set X is an edge and the factor-two bound is tight. For $d \geq 2$, the forbidden subcube contains a square, and Theorem 16 gives an infinite multiplicative gap. In particular, for the nonnegative objective $\Delta_F^a(x)$, the value over P_X is one, whereas the value over Q_X is zero whenever $d \geq 2$.

4.4 Induced stars

Suppose that X induces a star in the cube graph. Up to an automorphism of the cube, we may write $X = \{0, e^1, \ldots, e^r\}$ and $r = |X| - 1$. We assume $r \geq 2$. In these coordinates,

$$Q_X = \left\{ x \in [0,1]^n : \sum_{j=1}^n x_j \geq 1, \ \sum_{j \neq i} x_j \geq x_i \text{ for all } i \in [r] \right\}.$$

From [6], the exact convex-hull description consists of the cube bounds, the edge inequalities $\sum_{j \neq i} x_j \geq 1$ for $i \in [r]$, and the star inequality

$$\sum_{i=1}^r x_i + 2 \sum_{j=r+1}^n x_j \geq 2. \tag{1}$$

Since an induced star is square-free, Theorem 16 implies that Q_X has multiplicative gap at most two for all nonnegative affine objectives. The following proposition shows that the star inequality itself is violated by exactly this factor.

Proposition 19. *The maximum additive violation of* (1) *over* Q_X *is one, that is,*

$$\min\left\{\sum_{i=1}^{r} x_i + 2\sum_{j=r+1}^{n} x_j : x \in Q_X\right\} = 1.$$

Proof. Let $x \in Q_X$, and set $s := \sum_{j=1}^{n} x_j$ and $b := \sum_{j=r+1}^{n} x_j$. The left-hand side of (1) is $s + b$. The no-good cut associated with 0 gives $s \geq 1$, and $b \geq 0$. Hence the left-hand side is at least one.

The bound is attained by $\bar{x} = (e^1 + e^2)/2$. Indeed, $\sum_j \bar{x}_j = 1$ and $\sum_{j \neq i} \bar{x}_j \geq \bar{x}_i$ for every $i \in [r]$, so $\bar{x} \in Q_X$. The left-hand side of (1) at $\bar{x}$ is one. □

The factor-two objective gap is tight for proper stars. For example, take $c_1 = c_2 = 1$ and $c_j = M \geq 2$ for $j = 3, \ldots, n$. Then $\bar{x} = (e^1 + e^2)/2$ shows that $\min\{c^\top x : \; x \in Q_X\} = 1$, while every feasible binary point in P_X has cost at least two, and $e^1 + e^2$ is feasible with cost two. Hence $\min\{c^\top x : \; x \in P_X\} = 2$.

4.5 Tulip inequalities

We now consider tulip inequalities. For $\bar{x} \in \{0,1\}^n$, write $\bar{x}^i = \bar{x} \oplus e^i$ and $\bar{x}^{ij} = \bar{x} \oplus (e^i + e^j)$ for $i \neq j$. Assume that X contains a tulip rooted at $\bar{x}$. That is, for some distinct indices $i_1, \ldots, i_t$, with $t \geq 4$, the forbidden set contains

$$\bar{x}, \quad \bar{x}^{i_1}, \quad \bar{x}^{i_2}, \quad \bar{x}^{i_3}, \quad \bar{x}^{i_1 i_2}, \quad \bar{x}^{i_2 i_3}, \quad \bar{x}^{i_3 i_1}, \quad \bar{x}^{i_4}, \ldots, \bar{x}^{i_t}.$$

Let $I_T := \{i_1, \ldots, i_t\}$. For $j \in [n]$, define $y_j(x) := \bar{x}_j(1 - x_j) + (1 - \bar{x}_j)x_j$. The associated tulip inequality [6] is $L_T(x) \geq 3$, where

$$L_T(x) := \sum_{k=1}^{3} y_{i_k}(x) + 2\sum_{r=4}^{t} y_{i_r}(x) + 3\sum_{j \notin I_T} y_j(x).$$

Because a tulip contains the square $\bar{x}, \bar{x}^{i_1}, \bar{x}^{i_2}, \bar{x}^{i_1 i_2}$, Theorem 16 already implies that the no-good relaxation has infinite multiplicative gap for some nonnegative affine objective. The next result instead quantifies the violation of the tulip inequality itself.

Proposition 20. *For every tulip contained in* X,

$$\min\{L_T(x) : x \in Q_X\} = 1.$$

Proof. Since the root $\bar{x}$ belongs to X, the no-good cut for $\bar{x}$ gives $\sum_j y_j(x) = d_{\bar{x}}(x) \geq 1$ for every $x \in Q_X$. All coefficients in $L_T(x)$ are at least one, so $L_T(x) \geq d_{\bar{x}}(x) \geq 1$.

To prove tightness, define x^* by $x^*_{i_1} = x^*_{i_2} = 1/2$ and $x^*_j = \bar{x}_j$ for all $j \notin \{i_1, i_2\}$. Then $L_T(x^*) = 1$. Moreover, the coordinates i_1 and i_2 contribute exactly $1/2$ each to $d_v(x^*)$ for every $v \in \{0,1\}^n$. Hence $x^* \in Q_X$. □

4.6 Propeller inequalities

Finally, assume that the forbidden vertices induce a propeller. Up to an automorphism of the cube, we may write

$$X = \{0, e^1\} \cup \{e^i, e^1 + e^i : i = 2, \dots, t+1\}, \qquad t \geq 3.$$

The edge $0e^1$ is the axis of the propeller. The associated propeller inequality [6] is

$$\sum_{i=2}^{t+1} x_i + 2 \sum_{j=t+2}^{n} x_j \geq 2. \tag{2}$$

A propeller contains the square $0, e^1, e^2, e^1 + e^2$. Hence Theorem 16 implies that the objective gap of the no-good relaxation is infinite. The following proposition gives the exact violation of the propeller inequality.

Proposition 21. *The maximum additive violation of* (2) *over* Q_X *is* $3/2$*, that is,*

$$\min \left\{ \sum_{i=2}^{t+1} x_i + 2 \sum_{j=t+2}^{n} x_j : x \in Q_X \right\} = 1/2.$$

Proof. For $x \in Q_X$, define $a := x_1$, $A := \sum_{i=2}^{t+1} x_i$, and $B := \sum_{j=t+2}^{n} x_j$. The left-hand side of (2) is $\pi(x) := A + 2B$.

The no-good cut for 0 gives $a + A + B \geq 1$, while the no-good cut for e^1 gives $A + B \geq a$. Thus $2(A + B) \geq a + A + B \geq 1$, and therefore $A + B \geq 1/2$. Since $B \geq 0$, we get $\pi(x) = (A + B) + B \geq 1/2$.

The bound is attained by $\bar{x} = (e^1 + e^2)/2$. The coordinates 1 and 2 contribute exactly $1/2$ each to $d_v(\bar{x})$ for every $v \in \{0,1\}^n$, so $\bar{x} \in Q_X$. Moreover, $\pi(\bar{x}) = 1/2$. □

4.7 Gap after adding all studied inequalities

We close this section by considering the relaxation obtained after adding all inequalities studied above. Let R_X be the relaxation obtained by intersecting $[0,1]^n$ with every valid inequality of the following types whose forbidden configuration is contained in X: vertex inequalities, edge inequalities, subcube inequalities, star inequalities, tulip inequalities, and propeller inequalities. Here a p-dimensional subcube inequality is the inequality $\Delta_F^a(x) \geq 1$ associated with a face

$$C = \{y \in \{0,1\}^n : y_i = a_i \text{ for all } i \in F\} \subseteq X,$$

where $|F| = n - p$.

For $\ell \in \mathcal{L}_+$, define the objective gap of R_X with respect to P_X by

$$\mathrm{gap}_\ell(R_X, P_X) := \frac{\min\{\ell(x) : x \in P_X\}}{\min\{\ell(x) : x \in R_X\}},$$

whenever the denominator is positive. The worst-case gap is the supremum over all such objectives.

Theorem 22. *The relaxation* R_X *satisfies*

$$\min\{\ell(x) : x \in R_X\} \geq \frac{1}{n} \min\{\ell(x) : x \in P_X\}$$

for every $\ell \in \mathcal{L}_+$. *Hence* $\operatorname{gap}_\ell(R_X, P_X) \leq n$ *whenever the denominator is positive. Moreover, this bound is tight up to a constant factor: there exist connected sets* $X \subsetneq \{0,1\}^n$ *and nonnegative linear objectives* ℓ *for which* $\operatorname{gap}_\ell(R_X, P_X) = \Omega(n)$. *Therefore, the worst-case multiplicative gap after adding all studied inequalities is* $\Theta(n)$.

Proof. We first prove the upper bound. Let $\ell \in \mathcal{L}_+$. By complementing coordinates if necessary, which is an automorphism of the cube and preserves all inequality families under consideration, we may assume that

$$\ell(x) = \alpha + \sum_{i \in L} w_i x_i,$$

where $\alpha \geq 0$, $L \subseteq [n]$, and $w_i > 0$ for every $i \in L$. Let $g(x) := \sum_{i \in L} w_i x_i$ and $m := |L|$. If $m = 0$, then ℓ is constant and the claim is immediate. Hence assume $m \geq 1$.

Define

$$\beta := \min\{g(y) : y \in \{0,1\}^n \setminus X\}.$$

Then $\min\{\ell(x) : x \in P_X\} = \alpha + \beta$. If $\beta = 0$, then $\ell(x) \geq \alpha$ for every $x \in [0,1]^n$, and therefore $\min\{\ell(x) : x \in R_X\} \geq \alpha \geq (\alpha + \beta)/n$. Thus assume $\beta > 0$.

We claim that $g(x) \geq \beta/m$ for every $x \in R_X$. Suppose, for contradiction, that some $x \in R_X$ satisfies $g(x) < \beta/m$. Define $U := \{i \in L : x_i \geq 1/m\}$. Then

$$\sum_{i \in U} w_i \leq m \sum_{i \in U} w_i x_i \leq m g(x) < \beta.$$

Consider the subcube

$$C_U := \{y \in \{0,1\}^n : y_i = 0 \text{ for all } i \in L \setminus U\}.$$

Every $y \in C_U$ satisfies $g(y) \leq \sum_{i \in U} w_i < \beta$. By the definition of β, no vertex of C_U is allowed. Hence $C_U \subseteq X$.

Since $X \subsetneq \{0,1\}^n$, the subcube C_U is proper, and therefore $L \setminus U \neq \emptyset$. The subcube inequality associated with C_U belongs to R_X and is

$$\sum_{i \in L \setminus U} x_i \geq 1.$$

However, by the definition of U, we have $x_i < 1/m$ for all $i \in L \setminus U$, and therefore

$$\sum_{i \in L \setminus U} x_i < \frac{|L \setminus U|}{m} \leq 1,$$

contradicting $x \in R_X$. Thus $g(x) \geq \beta/m$ for every $x \in R_X$.

It follows that, for every $x \in R_X$,

$$\ell(x) = \alpha + g(x) \geq \alpha + \frac{\beta}{m} \geq \frac{\alpha + \beta}{m} \geq \frac{\alpha + \beta}{n}.$$

Taking the minimum over R_X gives

$$\min\{\ell(x) : x \in R_X\} \geq \frac{1}{n} \min\{\ell(x) : x \in P_X\}.$$

This proves the upper bound.

We now prove that the factor n cannot be improved to a constant. Fix $n \geq 5$ and set $r := \lfloor n/2 \rfloor$. Let X_r be the set of binary vectors with Hamming weight at most r. The graph induced by X_r is connected. Consider the objective $\ell(x) := \sum_{i=1}^{n} x_i$. Since the allowed binary vertices are exactly those with Hamming weight at least $r + 1$,

$$\min\{\ell(x) : x \in P_{X_r}\} = r + 1.$$

Let $\lambda := 2/(n-1)$ and $\bar{x} := \lambda \mathbf{1}$. We claim that $\bar{x} \in R_{X_r}$. Since $n \geq 5$, we have $\lambda \leq 1/2$. Hence, for every $v \in \{0,1\}^n$ and every coordinate i,

$$v_i(1 - \bar{x}_i) + (1 - v_i)\bar{x}_i \geq \lambda.$$

Thus, at $\bar{x}$, the left-hand side of any studied inequality is at least λ times the sum of its coefficients.

First consider a p-dimensional subcube $C \subseteq X_r$. Since every vertex of C has Hamming weight at most r, the dimension of C is at most r. The associated subcube inequality has coefficient sum $n - p$ and right-hand side one. Therefore its left-hand side at $\bar{x}$ is at least

$$\lambda(n - p) \geq \lambda(n - r) \geq 1.$$

Thus all subcube inequalities, including the vertex and edge inequalities, are satisfied.

Now consider a star inequality with t leaves. Its coefficient sum is $t + 2(n - t) = 2n - t$ and its right-hand side is two. Since $t \leq n$,

$$\lambda(2n - t) \geq \lambda n = \frac{2n}{n-1} \geq 2.$$

Thus every star inequality is satisfied.

Next consider a tulip inequality with parameter $t \geq 4$. Its coefficient sum is

$$3 + 2(t - 3) + 3(n - t) = 3n - t - 3,$$

and its right-hand side is three. Since $t \leq n$,

$$\lambda(3n - t - 3) \geq \lambda(2n - 3) = \frac{2(2n-3)}{n-1} \geq 3$$

for every $n \geq 5$. Thus every tulip inequality is satisfied.

Finally consider a propeller inequality with parameter $t \geq 3$. Its coefficient sum is

$$t + 2(n - t - 1) = 2n - t - 2,$$

and its right-hand side is two. Since the propeller uses $t + 1$ distinct coordinate directions, $t \leq n - 1$. Therefore

$$\lambda(2n - t - 2) \geq \lambda(n - 1) = 2.$$

Thus every propeller inequality is satisfied. We have proved that $\bar{x} \in R_{X_r}$.

Hence,

$$\min\{\ell(x) : x \in R_{X_r}\} \leq \ell(\bar{x}) = n\lambda = \frac{2n}{n-1}.$$

Therefore

$$\operatorname{gap}_\ell(R_{X_r}, P_{X_r}) \geq \frac{r+1}{2n/(n-1)} = \frac{(r+1)(n-1)}{2n} = \Omega(n).$$

Combining this lower bound with the general upper bound proves that the worst-case multiplicative gap is $\Theta(n)$. □

5 Concluding remarks

We studied the convex hull obtained from the binary cube after excluding a set of forbidden vertices, with emphasis on the relation between the number of forbidden vertices, the number of facets, extension complexity, and the strength of no-good cut relaxations. The results show a sharp qualitative contrast between fixed, polynomial, and unrestricted forbidden sets. When the number of forbidden vertices is fixed, both the number of facets and the extension complexity remain linear in the ambient dimension, up to constants depending only on the number of forbidden vertices. In contrast, polynomially many forbidden vertices can already force superpolynomially many facets, while the unrestricted regime recovers the extremal behavior of general 0/1 polytopes.

For extension complexity, the general formulation gives an upper bound of order $n|X|$, while the lower-bound constructions show that polynomially many forbidden vertices can force extension complexity at least $|X|^{1/2-o(1)}$ along suitable sequences. This leaves a substantial gap between the best known lower and upper bounds as functions of both n and $|X|$. Closing this gap, or identifying structural conditions on X under which the upper bound is essentially tight, appears to be a natural direction for further work.

The gap results show that no-good cuts and several known families of valid inequalities may still leave large objective gaps. Thus, although no-good cuts are a convenient and broadly applicable modeling device, their polyhedral strength depends strongly on the geometry induced by the forbidden set. A more refined understanding of this geometry could lead to sharper descriptions for special classes of forbidden configurations and to stronger cut families for binary optimization models in which infeasible assignments are generated or discovered dynamically.

References

[1] Gustavo Angulo, Shabbir Ahmed, Santanu S. Dey, and Volker Kaibel. Forbidden vertices. *Mathematics of Operations Research*, 40(2):350–360, 2015.

[2] Imre Bárány and Attila Pór. On 0-1 polytopes with many facets. *Advances in Mathematics*, 161(2):209–228, 2001.

[3] Yohann Benchetrit, Samuel Fiorini, Tony Huynh, and Stefan Weltge. Characterizing polytopes in the 0/1-cube with bounded Chvátal-Gomory rank. *Mathematics of Operations Research*, 43(3):718–725, 2018.

[4] J. F. Benders. Partitioning procedures for solving mixed-variables programming problems. *Numerische Mathematik*, 4:238–252, 1962.

[5] Gianni Codato and Matteo Fischetti. Combinatorial Benders' cuts for mixed-integer linear programming. *Operations Research*, 54(4):756–766, 2006.

[6] Gérard Cornuéjols and Dabeen Lee. On some polytopes contained in the 0, 1 hypercube that have a small Chvátal rank. *Mathematical Programming*, 172(1):467–503, 2018.

[7] Özgün Elçi and John Hooker. Stochastic planning and scheduling with logic-based Benders decomposition. *INFORMS Journal on Computing*, 34(5):2428–2442, 2022.

[8] Samuel Fiorini, Volker Kaibel, Kanstantsin Pashkovich, and Dirk Oliver Theis. Combinatorial bounds on nonnegative rank and extended formulations. *Discrete Mathematics*, 313(1):67–83, 2013.

[9] Tamás Fleiner, Volker Kaibel, and Günter Rote. Upper bounds on the maximal number of facets of 0/1-polytopes. *European Journal of Combinatorics*, 21(1):121–130, 2000.

[10] Dionysios Gatzouras, Apostolos Giannopoulos, and Niki Markoulakis. On the maximal number of facets of 0/1 polytopes. In *Geometric Aspects of Functional Analysis: Israel Seminar 2004–2005*, pages 117–125. Springer, 2007.

[11] John N Hooker. Logic-based Benders decomposition for large-scale optimization. In *Large scale optimization in supply chains and smart manufacturing: Theory and applications*, pages 1–26. Springer, 2019.

[12] Mojtaba Hosseini and John Turner. Deepest cuts for Benders decomposition. *Operations Research*, 73(5):2591–2609, 2025.

[13] Pavel Hrubeš and Navid Talebanfard. On the extension complexity of polytopes separating subsets of the boolean cube. *Discrete & Computational Geometry*, 70(1):268–278, 2023.

[14] Robert G Jeroslow. On defining sets of vertices of the hypercube by linear inequalities. *Discrete Mathematics*, 11(2):119–124, 1975.

[15] Volker Kaibel and Stefan Weltge. Lower bounds on the sizes of integer programs without additional variables. *Mathematical Programming*, 154(1):407–425, 2015.

[16] Gilbert Laporte and François V. Louveaux. The integer L-shaped method for stochastic integer programs with complete recourse. *Operations Research Letters*, 13(3):133–142, 1993.

[17] Thomas Rothvoß. Some 0/1 polytopes need exponential size extended formulations. *Mathematical Programming*, 142(1–2):255–268, 2013.

[18] Mihalis Yannakakis. Expressing combinatorial optimization problems by linear programs. In *Proceedings of the twentieth annual ACM symposium on Theory of computing*, pages 223–228, 1988.

[19] Günter M Ziegler. *Lectures on polytopes*, volume 152. Springer Science & Business Media, 2012.